

AD-A196 142

ref [3]
DTIC FILE COPY

2

STATISTICAL IMAGE RESTORATION AND REFINEMENT

Christopher Jennison and Michael Jubb

School of Mathematical Sciences

University of Bath, BATH BA2 7AY, U.K.

DTIC
ELECTE
JUN 07 1988
S D

Summary

We consider the problem of reconstructing an image from a noisy record. We describe existing methods due to Geman and Geman (1984) and Besag (1986) which use a Markov random field model for the true scene but assume that each pixel consists of a single colour. In order to improve the quality of the restoration at the boundary of regions of different colours we extend these methods to allow pixels to contain two regions of colour separated by a single straight line. An algorithm for performing the reconstruction is presented and illustrated by an example.

1. Introduction

We consider a rectangular region partitioned into pixels labelled $1, 2, \dots, n$. Each pixel is coloured black or white and the colour of pixel i is denoted by x_i which takes the value 0 for white and 1 for black. The x_i are unobserved. We work instead from the observed record y_i which consists of x_i plus added noise. We denote the whole scene by $x = \{x_i; i=1, \dots, n\}$ and the set of records by $y = \{y_i; i=1, \dots, n\}$. The noise distribution will be assumed to be known but if this were not the case, it could be established by studying training data.

Recent developments in statistical restoration methods use a Bayesian approach. The *maximum a posteriori* (MAP) estimate of the true scene is the value of x which maximises $P(x|y)$, the conditional probability of x given the record y . By Bayes' theorem

$$P(x|y) \propto l(y|x) p(x), \quad (1.1)$$

where $l(y|x)$ is the conditional likelihood of the observed record, y , given the true colouring, x , and $p(x)$ is the prior probability of x .

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

88 5 17 045

172-2911

9



Abstract

←



✓	
□	
□	

per NP

0712
0085
IN 00000000

A-1

$$\frac{1}{2\sigma^2} (y_i - x_i)^2 + \beta Z(x_i)$$

where $Z(x_i)$ is the number of neighbours of pixel i in the current restoration which are of the opposite colour to x_i . The method proceeds by scanning the scene, successively updating each pixel until convergence is reached. This will normally occur at a local rather than global maximum of $P(x|y)$, but, given the possibility of undesirable long range dependencies in the MRF model, this is not a serious drawback and might even be an advantage.

2. Split Pixels

So far we have considered scenes in which each pixel is coloured wholly one colour. We now allow pixels in the true scene to be coloured partly black and partly white. Each record y_i is distributed as Gaussian with variance σ^2 and mean p_i , the proportion of pixel i which is coloured black. The restoration methods that we have previously discussed can be used for this problem by proceeding as if the pixels were only of one colour but the quality of the restoration at the edges of objects or regions will obviously be poor. Instead, we can allow pixels in the restored image to be coloured partly black and partly white. The simplest form of this is to quarter each pixel and allow it to be filled with the most likely of the 2^4 configurations. This method, proposed by Jennison (1986) uses a modified version of ICM, firstly iterating at full pixel size and subsequently restoring the quarters; in the second stage the same form of MRF model is used for the *subpixels* as is originally used for full pixels. This method appears to work well and has prompted work into the further breakdown of pixels.

For further refinement we can either (i) consider an $m \times m$ breakdown of each pixel or (ii) use continuous lines within the pixel to represent the edge. The implementation of (i) requires the minimisation of

$$\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - \frac{1}{m^2} \sum_{j=1}^m \sum_{k=1}^m x_{ijk})^2 + \beta \sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^m Z(x_{ijk}),$$

where the subscript ijk refers to subpixel j, k within pixel i ; x_{ijk} takes value 0 or 1 and $Z(x_{ijk})$ is the number of subpixel neighbours of subpixel ijk in the current restoration which are of the opposite colour to x_{ijk} . Note that subpixels at the edge of a pixel will have some subpixel neighbours contained in an adjacent pixel. We can see that as m increases this minimisation becomes computationally cumbersome. Also, it offers only an approximation to (ii) and it turns out to be easier to pass to the limit and work directly with continuous solutions.

The most basic form of (ii) allows a single straight line edge within each pixel and it is the implementation of this that we shall describe. It is no longer meaningful to talk of discrepant pixel or subpixel pairs and we replace the second term of (1.2) by a multiple of the total length of edge in the reconstruction x . Thus, the restored image is chosen to minimise

$$\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - p_i(x))^2 + \beta' L(x), \quad (2.1)$$

over images x made up of pixels x_i , $i=1, \dots, n$ either of a single colour or divided into two regions of different colours by a single straight line; $p_i(x)$ denotes the proportion of black in pixel i ; $L(x)$ is the total edge length in scene x and β' is a fixed constant related to the β used earlier.

An advantage of edge length as a measure is that the penalty is rotationally invariant, i.e. remains constant throughout all rotations of the scene within the region. This could not be obtained using discrepant pairs as a measure although it has been shown by our colleague Robin Sibson that this variability can be minimised using a down weighting of $1/\sqrt{2}$ for the diagonal adjacencies.

3. The Restoration Algorithm

The restoration is done in three stages, the first two of which have already been described :

Stage 1 : ICM to convergence on full size pixel grid.

Stage 2 : ICM to convergence on 2×2 pixel grid.

Stage 3 : Updating process on the line segments representing the edges.

Stage 3 requires that we now regard the reconstruction as a series of line segments separating the two colours. An initial representation is obtained in a straightforward way from the end product of Stage 2. The updating process treats pixels in pairs, selecting the best place for two edges to meet, given the current restoration of neighbouring pixels.

As an example, consider the configuration at pixels i and j shown in Figure 1. The distances a and b are determined by the current colouring of neighbouring pixels and treated as constant for the moment. The distance W is chosen to minimise the contribution from pixels i and j to the total penalty (2.1), i.e.

$$Pen(W) = \frac{1}{2\sigma^2} \sum_{k=i,j} (y_k - p_{kW})^2 + \beta'(e_{iW} + e_{jW}), \quad (3.1)$$

where e_{kW} is the length of edge in pixel k when the join is at W and p_{kW} is the proportion of black in pixel k when the join is at W .

For the case shown in Figure 1, this penalty is

$$g(W) = \frac{1}{2\sigma^2} \{ (y_i - a - \frac{1}{2}(W-a))^2 + (y_j - b - \frac{1}{2}(W-b))^2 \} \\ + \beta' \{ \sqrt{1+(W-a)^2} + \sqrt{1+(W-b)^2} \}.$$




Fig. 1. Updating the position of edges in pixels i and j .

This can not be minimised directly but the form of

$$\frac{dg(W)}{dW} = \frac{1}{4\sigma^2} (2W + a - 2y_i + b - 2y_j) + \beta' \left[\frac{(W-a)}{\sqrt{1+(W-a)^2}} + \frac{(W-b)}{\sqrt{1+(W-b)^2}} \right]$$

suggests an iterative approach. Given an approximate solution W_{s-1} we solve

$$\frac{1}{4\sigma^2} (2W_s + a - 2y_i + b - 2y_j) + \beta' \left[\frac{(W_s-a)}{\sqrt{1+(W_{s-1}-a)^2}} + \frac{(W_s-b)}{\sqrt{1+(W_{s-1}-b)^2}} \right] = 0$$

to obtain

$$W_s = \frac{4\sigma^2\beta' \left[\frac{a}{\sqrt{1+(W_{s-1}-a)^2}} + \frac{b}{\sqrt{1+(W_{s-1}-b)^2}} \right] + (2y_i - a + 2y_j - b)}{2 + 4\sigma^2\beta' \left[\frac{1}{\sqrt{1+(W_{s-1}-a)^2}} + \frac{1}{\sqrt{1+(W_{s-1}-b)^2}} \right]}$$

Starting from any sensible initial value, W_0 , accuracy to 3 decimal places was achieved after at most four iterations. In practice we take W_0 to be the value of W prior to this update.

Different forms of (3.1) are possible depending on which neighbours of pixels i and j contain both colours. There are only four distinct cases that may arise and these are shown in Figure 2.




Fig. 2. Possible configurations of edges in two neighbouring pixels.

We have shown the method of solution for case (i) and cases (ii) - (iv) are solved in a similar way. All other cases can be reduced to one of the above by means of exchanging and/or inverting the pixels and their colours.

The most natural order of updating the edge pixels would seem to be to follow an edge around, updating each join in turn, completing circuits of the edge until convergence. An alternative method is to update every k^{th} join around the circuit, therefore completing k laps before each pixel has been updated once. Initial results suggest that this provides additional stability in the updating process; we have found the value $k = 3$ to give particularly good results.

4. An example

We illustrate the methods we have described with an artificial example. Figure 3a shows a true image and the superimposed pixel grid. The record from which a restored image was constructed obtained by generating a Gaussian random variable for each pixel with mean equal to the proportion of the pixel coloured black in the true image and variance 0.1^2 . Figure 3b is the reconstruction after stage 1, in which the ICM method with $\beta=1$ has been used, treating each pixel as either completely black or completely white. Note that this is a rather poor approximation to the true image but it is the best that can be done without dividing pixels. Subdividing each pixel into four in stage 2 produces the reconstruction in Figure 3c: the amounts of black in each full pixel are now much closer to the corresponding records and the divisions of split pixels match up well with the true image. Proceeding to stage 3, using $\beta'=2$, gives the final reconstruction shown in Figure 3d, despite the coarseness of the original pixel grid and the addition of noise to the record, this reconstruction is barely distinguishable from the true image.




Fig 3a True image




Fig 3b Reconstruction after stage 1




Fig 3c Reconstruction after stage 2




Fig 3d Final reconstruction

5. Further extensions

(a) Consider a pixel which has true colouring as shown in Figure 4. Clearly the straight line approximation to this edge will be poor and could have an adverse effect on the reconstruction of neighbouring pixels and pixels further along the edge. This may be overcome using a more intricate restoration method, e.g. allowing two straight lines meeting at some point within a pixel.




Fig. 4. A pixel containing a boundary that can not be approximated well by a single straight line.

(b) The method presented in this paper can be extended to scenes containing more than two different colours. Where any two regions meet we can adjust the algorithm to provide a continuous line join. More computation is required to find the best colouring for a pixel in which three or more regions meet.

References

- Besag, J. E. (1974) Spatial interaction and the statistical analysis of lattice systems. *J. Royal Statist. Soc. B*, 36, 192-236.
- Besag, J. E. (1986) On the statistical analysis of dirty pictures. *J. Royal Statist. Soc. B*, 148, 259-302.
- Geman, S. and Geman, D. (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, 6, 721-741.
- Greig, D. M., Porteous, B.T. and Seheult, A.H. (1986) Contribution to the discussion of Besag (1986) *J. Royal Statist. Soc. B*, 148, 282-284.
- Jennison, C. (1986) Contribution to discussion of Besag (1986) *J. Royal Statist. Soc. B*, 148, 288-289.